

ACTS – A COLLECTION OF HIGH PERFORMING SOFTWARE TOOLS FOR SCIENTIFIC COMPUTING

L.A.DRUMMOND AND O.A. MARQUES

*Lawrence Berkeley National Laboratory, One Cyclotron Road, 50F MS-1650
Berkeley, California, USA*

E-mail: LADrummond@lbl.gov, and OAMarques@lbl.gov

During the past decades there has been a continuous growth in the number of physical and societal problems that have been successfully studied and solved by means of computational modeling and simulation. Further, many new discoveries depend on high performance computer simulations to satisfy their demands for large computational resources and short response time. The Advanced Computational Software (ACTS) Collection brings together a number of general-purpose computational tool development projects funded and supported by the U.S. Department of Energy (DOE). These tools make it easier for scientific code developers to write high performance applications for parallel computers. They tackle a number of computational issues that are common to a large number of scientific applications, mainly implementation of numerical algorithms, and support for code development, execution and optimization. The ACTS collection promotes code portability, reusability, reduction of duplicate efforts, and tool maturity. This paper presents a brief introduction to the functionality available in ACTS. It also highlights the tools that are in demand by Climate and Weather modelers.

1 INTRODUCTION

The United States Department of Energy's (DOE) Advanced Computational Software (ACTS) Collection is a set of computational tools developed primarily at DOE laboratories and is aimed at simplifying the solution of common and important computational problems. These tools are freely available with minor licensing agreements for their use and distribution. The use of the tools reduces the development time for new codes and the tools provide functionality that might not otherwise be available. All this potential cannot be achieved, however, if the tools are not used effectively or not used at all.

The ACTS Project brings together software development from DOE laboratories, and in some cases in collaboration with universities, to the hands of a scientific community that is broader than the community in a single DOE laboratory. Consequently, ACTS has promoted reusability over duplication of efforts and tool interoperability over "hard-wired" and intrusive library interfaces that only serve a specific project.

LADrummond-2002 submitted to World Scientific 2/3/03 : 10:06 AM
--

1/15

ACTS Tools are mostly library with interfaces to programs written in Fortran, C and C++. They are primarily designed to run on high performance computing architectures. Portability and performance are both considerations in their design and adaptability to emerging computer technologies. At the National Energy Research Scientific Computing Center (NERSC), we are currently providing a comprehensive support to the users of tools under ACTS. This support includes training and education, help prototyping scientific and engineering codes, and independent software evaluation of the tools.

This paper presents a short introduction to the ACTS Collection. For more materials, tutorials and references visit the ACTS Information Center¹. In Section 2, we present a list of the tools currently available in the ACTS Collection. Conclusions and future plans are presented in Section 3.

2 WHAT CAN ACTS DO FOR YOU?

Generally climate and weather models utilize several numerical algorithms that may involve the iterative solution of linear or non-linear system equations, computations of a fraction or full spectrum of eigenvalues and eigenvectors associated with the physical phenomena, an optimization or minimization of a function, filtering techniques, and combinations of all of these. As the resolution and domain of these models is increased demand for computational resources is also increased in a non-linear manner.⁴⁻¹⁷

Therefore, these models should be developed in a way that they not only optimized the use of computational resources but also can easily adapt to changes in the physical parameterizations, and computational environments. Currently, there are 7 libraries in the ACTS Collection that efficiently implement a wide variety of numerical algorithms (see Section 2.1) that scale very well in a number of computer platform.

The ACTS Collection also provides a number of other useful services for large-scale numerical simulations found in climate and weather applications. These services include support for model coupling, data manipulation, remote computational steering, code interoperability support, and grid management and generation.

The libraries currently available in the ACTS Collection² can be categorized by their functionality in 5 different groups, Numerical Tools, Tools That Support Code Development, Tools That Support

Code Execution, Tools That Support Library Development, and Interoperability.

2.1 Numerical Tools

Aztec (developed at Sandia National Laboratories, SNL^{2,25,28,29}) is a library that provides algorithms for the iterative solution of large sparse linear systems arising in scientific and engineering applications. It is a stand-alone package comprising a set of iterative solvers, preconditioners and matrix-vector multiplication routines. Users are not required to provide their own matrix-vector multiplication routines or preconditioners in order to solve a linear system. The Aztec library is written in C and is also callable from Fortran. It is portable to most parallel platforms since it uses MPI to perform data communication. Overall, the package was designed to be easy to use. The user may input the linear system in a simple format and Aztec will perform the necessary transformations for the matrix-vector multiplication and preconditioning. After the transformations, the iterative solvers can run efficiently. If the input matrix is suitably partitioned, the efficiency can be further enhanced.

Aztec has been superseded by the Trilinos solver framework²⁸. Aztec functionality is available in Trilinos through an Object-Oriented interface named AztecOO. Former users of Aztec continue to have access to its functionality.

Hypre^{2,8} (developed at Lawrence Livermore National Lab, LLNL) is a library for solving large, sparse linear systems of equations on massively parallel computers. The main features of this library are: scalable preconditioners, implementation of a suit of common iterative methods (these include Conjugate Gradient and GMRES for symmetric and unsymmetric matrices, respectively), intuitive grid-centric interfaces, and dynamic configuration of parameters. Hypre works for users with different levels of expertise and has user-defined interfaces for multiple languages.

OPT++² (developed at SNL) is an object-oriented nonlinear optimization package for serial architectures. It solves optimization problems of the form

$$\min_{x \in R^n} f(x), h_i(x) = 0, i = 1, \dots, p, g_j(x) \geq 0, j = 1, \dots, m,$$

in which the user specifies the function f (and, when available, its first and second analytical derivatives) and the functions h and g . OPT++ provides four solution algorithms: a Newton method, a finite-difference Newton method, a Quasi-Newton method, and a nonlinear conjugate gradient method.

PETSc^{2,4,5,21}, the Portable, Extensible Toolkit for Scientific computation (developed at Argonne National Laboratory, ANL) provides sets of tools for the parallel (as well as serial,) numerical solution of PDEs that require solving large-scale, sparse linear and nonlinear systems of equations. PETSc includes nonlinear and linear equation solvers that employ a variety of Newton techniques and Krylov subspace methods. PETSc provides several parallel sparse matrix formats, including compressed row, block compressed row, and block diagonal storage.

PETSc is one of the most popular and matured tools currently available in the ACTS Toolkit. PETSc has a growing number of application users and development projects that re-use some of the functionality available in PETSc (like TAO and SLEPc²⁷). The PETSc team has long worked on the development of interoperable interfaces with other high performing tools. PETSc has also been used in a wide range of scientific applications involving, among others, the complex task of optimally porting legacy code to high performing computers^{12, 14}.

PVODE² (developed at LLNL) actually refers to a trio of closely related solvers: PVODE, for systems of ordinary differential equations, KINSOL, for systems of nonlinear algebraic equations, and IDA, for systems of differential-algebraic equations. These solvers have some modules in common, primarily a module of vector kernels, and a generic linear system solver based on a scaled preconditioned GMRES method. PVODE is a solver for large systems of ordinary differential equations on parallel machines. It contains methods for the solution of both stiff and non-stiff initial value problems. Integration methods include the variable coefficient forms of the Adams and backward differentiation formula methods. The linear systems that must be solved during the implicit time stepping are solved with iterative, preconditioned Krylov solvers. The user can either supply a preconditioner or use one that is included in the PVODE package. PVODE is an extension of the

sequential package known as CVODE, which has been widely distributed and used²⁰.

The PVOE trio is most useful for solving large differential and nonlinear algebraic systems that arise in a variety of applications. Important DOE applications include chemical kinetics, atmospheric chemistry, semiconductors, and structural or mechanical systems.

ScaLAPACK^{2,3}(developed at the University of Tennessee, Knoxville; University of California, Berkeley; and ORNL) is a library of high performance linear algebra routines for distributed-memory message-passing MIMD computers and networks of workstations supporting PVM or MPI. It is a continuation of the LAPACK project¹⁵, which designed and produced analogous software for workstations, vector supercomputers, and shared-memory parallel computers. The goals of both projects are efficiency (to run as fast as possible), scalability (as the problem size and number of processors grow), reliability (including error bounds), portability (across all important parallel machines), flexibility (so users can construct new routines from well-designed parts), and ease of use (by making the interface to LAPACK and ScaLAPACK look as similar as possible). The ScaLAPACK library contains routines for solving systems of linear equations, least squares, eigenvalue and singular value problems. They can also handle many associated computations such as matrix factorizations or estimation of condition numbers.

ScaLAPACK is intended for use in large-scale applications that require numerical manipulation of large dense or band matrices. Within NERSC's user community, we have seen various routines from ScaLAPACK being used. The applications using ScaLAPACK include material sciences, computational chemistry and astrophysics.

SuperLU² (developed at the University of California, Berkeley; and NERSC) is a general-purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations on high performance machines. The library is written in C and is callable from either C or Fortran. The library routines perform an LU decomposition with numerical pivoting and triangular system solves through forward and back substitution. The LU factorization routines can handle non-square matrices but the triangular solves are performed only for square matrices. The matrix columns may be preordered (before factorization) either through library or user-supplied routines. This reordering for sparsity is completely separate from the factorization. Working

precision iterative refinement subroutines are provided for improved backward stability. Routines are also provided to equilibrate the system, estimate the condition number, calculate the relative backward error, and estimate error bounds for the refined solutions.

TAO², the Toolkit for Advanced Optimization (developed at ANL), focuses on large-scale optimization software, including nonlinear least squares, unconstrained minimization, bound constrained optimization, and general nonlinear optimization. There are a variety of software tools for solving the aforementioned problems; however, only TAO offers an Object-Oriented solution that provides a flexible optimization toolkit capable of addressing issues of portability, versatility and scalability in many computational environments. The algorithms in TAO place strong emphasis on the reuse of external tools where appropriate. TAO's design enables bidirectional connection to lower-level linear algebra support (such as parallel sparse matrix data structures) that is available in toolkits like PETSc.

2.2 Tools that Support Code Development

Global Arrays^{2,18} (GA, developed at Pacific Northwest National Laboratory, PNNL) is a library for writing parallel programs that use large arrays distributed across processing nodes and that offers a shared-memory view of distributed arrays without destroying their NUMA characteristics. The library has both Fortran and C interfaces. Originally developed to support arrays as vectors and matrices (one or two dimensions), it currently supports up to seven dimensions in Fortran and even more in C. GA offers two types of operations: collective operations (require participation and synchronization of all processes) and local operations (may be invoked independently by all processes). GA also comes with a visualizer that uses trace files to animate array access patterns. Its main purpose is to analyze the impact of distribution on performance.

GA is meant to complement rather than replace the message-passing model, and is interoperable with MPI. GA was originally developed for use in the NWChem computational chemistry package¹⁹, but has also been used in other chemistry packages.

Overture² (developed at LLNL) is a set of Object-Oriented tools for solving computational fluid dynamics and combustion problems in complex moving geometries. It has been designed for solving problems

on a structured grid or a collection of structured grids. It can use curvilinear grids, adaptive mesh refinement, and the composite overlapping grid method to represent problems involving complex domains with moving components. Overture programs are written at a very high level, using data-parallel array expressions in the style of High Performance Fortran (HPF). They can achieve high performance (comparable to Fortran) thanks to a preprocessor (C++ to C++) called ROSE. Effectively, ROSE is a replacement for the expression template technique of POOMA. Overture has aggregate array operations and tightly integrated graphical features based on OpenGL. AMR++, a package that directly supports adaptive mesh refinement methods, is built on top of Overture.

In October 2001 the Overture Project was named one of the 101 most important discoveries supported by the DOE Office of Science in the past 25 years.

2.3 Tools that Support Code Execution and Performance Monitoring

CUMULVS^{2,13} (Collaborative User Migration, User Library for Visualization and Steering, developed at Oak Ridge National Laboratory, ORNL) is a software framework that enables programmers to incorporate fault-tolerance, interactive visualization and computational steering into existing parallel programs. The CUMULVS software consists of two libraries, one for the application program, and one for the visualization and steering front-end (called the "viewer"). CUMULVS handles collecting and transferring distributed data fields to the viewers and oversees adjustments to steering parameters in the application. It also manages the dynamic attachment and detachment of multiple independent viewers to a running parallel application. In addition, CUMULVS provides a user-directed checkpoint/restart mechanism to enable users to integrate fault tolerance to a running parallel application.

Among others, CUMULVS is very effective for monitoring and steering remote parallel execution and allowing the construction of attractive user interfaces that have greatly helped users that are not familiar with all the particulars of a given code and/or the complexities involved in parallel computing.

Globus^{2,10,11} provides a bag of services for the creation of computational Grids and tools with which applications can be developed to access the Grid. Globus itself provides the following core services: communication services, resource allocation and process management. It also provides a standard interface to local resource management systems, allowing computational Grid tools and applications to express resource allocation and process management requests in terms of a standard application programming interface (API) and protocol, without constraining individual sites to any specific resource management tool. A standard protocol (LDAP, the Lightweight Directory Access Protocol) allows for discovering, publishing, and accessing information about the configuration and status of the computational Grid.

Globus developers were awarded the 1997 Global Information Infrastructure (GII) Next Generation Award for their work in advancing the technology and application of high performance distributed computing. The Globus technology is now used worldwide and it has provided a leading infrastructure for the deployment of and collaborations around Grid technology. There are three main kinds of users of this technology: the developers of Grid services, Grid-enabled application developers and users of the applications. Nowadays, all three kinds of users have to rely on their institutional agreement and provisions to work with the Globus technology, therefore the current support to Globus users requires a larger orchestration of personnel and resources than any of the other ACTS tools.

*PAWS*² (Parallel Application Work Space, developed at Los Alamos National Laboratory, LANL) provides a framework for coupling parallel applications within a component-like model. Central to the design of PAWS is the coupling of parallel applications using parallel communication channels. The coupled applications can be run on different machines, and the data structures in each coupled component can have different parallel distributions. PAWS is able to carry out the communication without having to resort to global gather/scatter operations. Instead, point-to-point transfers are performed in which each node sends segments of data to remote nodes directly and in parallel.

With the abrupt urge to couple multi-disciplinary codes there are high and increasing demands for tools like PAWS. The tool is currently under consideration by several research groups and we foresee the integration of functionalities available in other coupling efforts and PAWS through the use of component technology.

TAU^{2,26} (Tuning and Analysis Utilities, developed at the University of Oregon and LANL) is a set of tools for analyzing the performance of C, C++, Fortran and Java programs. It is normally used in three steps:

1. Instrument the program by inserting TAU macros into the program (this can be done automatically for C++ programs)
2. Run the program, a trace file is then automatically generated.
3. View the trace file with a TAU visualizer (RACY) or a third-party visualizer (such as VAMPIR).

TAU collects much more information than what is available through *prof* or *gprof*, the standard Unix utilities. Also available through TAU are: per-process, per-thread and per-host information (supports *pthreads*), inclusive and exclusive function times, profiling groups that allow you to organize data collection, access to hardware counters on some systems, per-class and per-instance information, separate data for each template instantiation, start/stop timers for profiling arbitrary sections of code, and support for collection of statistics on user-defined events.

TAU can be used by anyone writing a C, C++, Fortran or Java application and that wants to understand where the performance bottlenecks are. The hurdles of inserting profiling functions is minimized by two factors: only a single statement is needed to profile a function, and profiling can be disabled completely so that profiling insertions don't need to be taken out. Users of TAU have been enthusiastic about its usefulness.

2.4 Tools that Support Library Development

ATLAS² (Automatically Tuned Linear Algebra Software, developed at the University of Tennessee, Knoxville) and **PHiPAC**² (Portable High Performance ANSIC, developed at the University of California, Berkeley, and the International Computer Science Institute) are tools for the automatic generation of optimized numerical software for modern computer architectures and compilers. These tools have both initially focused on level-three BLAS operations (matrix-matrix multiplications) and also a few routines from LAPACK, which have high potential for optimization. Traditionally, the optimization of these routines has been a tedious, architecture-dependent, hand coding process. Codes automatically generated by ATLAS and PHiPAC have been able to meet and even exceed the performance of the vendor-supplied, hand-optimized BLAS on a range of platforms. ATLAS and PHiPAC achieve

performance through loop unrolling, explicit removal of unnecessary dependencies in code blocks, and use of machine-sympathetic C constructs. Code generators are parameterized and scripts are used to find the optimal choice of parameters for a given architecture and compiler. Research continues on these tools to extend them to new BLAS 3 operations and other basic linear algebra calculations, including sparse operations.

ATLAS and PHiPAC are of primary interest to computer architects and administrators and are not tools that the application programmer would use, although the technology that they are enabling may well eventually spill over into tools for application programmers.

2.5 Tools that Support Interoperability

The Objective of the Common Component Architecture Forum^{6,7} (CCA Forum) is to define a minimal set of standard features that a High-Performance Component Framework has to provide, or can expect, in order to be able to use components developed within different frameworks. Such standard will promote interoperability between components developed by different teams across different institutions. The Center for Component Technology for Terascale Simulation Software (CCTSS) is dedicated to the development of a component-based software development model suitable for the needs of high-performance scientific simulation, particularly the CCA. The Center is funded by the U. S. Dept. of Energy (DOE) as an Integrated Software Infrastructure Center (ISIC) under the Scientific Discovery through Advanced Computing (SciDAC²⁴) program and includes members from Argonne, Livermore, Los Alamos, Oak Ridge, Pacific Northwest, and Sandia National Laboratories, Indiana University and the University of Utah.

3 CONCLUSIONS AND FUTURE PLANS.

For the past 15 years or so, there have many discoveries in climate related research and improvements in weather prediction that are due in part to better computational software tools and resources. However, there is still a gap between tool developers and application developers which leads to duplication of efforts. Without projects like ACTS, application developers will continue to design and implement codes using techniques that are already available from other sources. Quite

often these new developments are far from optimal because of the application developers' inexperience with all the different issues that lead to optimal performance. In many cases, application developers only consult sources like Numerical Recipes²² that do not address platform optimization and parallelism, algorithm robustness, and language specific optimization issues²³.

One of the main concerns that ACTS users have expressed to us is the longevity of support from tool developers and required evolution of the software as the hardware technology continues to evolve and the complexity of the scientific application continues to grow. Inasmuch as a solid base of reusable tools is utilized inside newer tool developments, users are guaranteed the evolution of the tools; thus long-term support and functionality are also guaranteed.

The tools in the ACTS Collection are freely available to users in the United States, as well as at the international level. Applications and users play an important role in hardening tools. The main parameters for maturity are portability, robustness, acceptance, and long-term support. It is particularly the interactions with real users and real applications that have made the software mature, portable, robust and better documented. In turn, mature software will be widely accepted inside a given scientific community. The current framework of the ACTS Collection has promoted the tools to a wider national and international audience, thus increasing not only the visibility of the tools worldwide but also the range of users and applications.

There is a demand for tool interoperability and more uniformity in the documentation and user interfaces. Users want to experiment with functionalities available in a subset of the ACTS tools, and finding similar user interfaces and comparable levels of support and documentation makes this task even simpler and risk free. Furthermore, the computational challenges at hand demand new software developments that interact with legacy code practices, data and computer languages. ACTS provides a natural infrastructure to put in practice all the code, data and language interoperability required by these new challenges.

In our computational science community there is a need for an intelligent and dynamic catalog/repository of high performance tools. The need for a centralized software and reference repository is vital for preventing the duplication of efforts. Currently, the ACTS Information Center provides pointers to tools currently funded under ACTS with the accumulated expertise from tool users and scientific domains that are served by the tools. Additionally, we envision the inclusion of pointers

to tools offering functionality not currently available in the ACTS Collection as well as fair comparisons with other tools that offer functionality that overlaps with the ACTS tools.

Acknowledgements

The ACTS Collection project is funded by the US Department of Energy, office of Mathematical, Information and Computation Science.

References

1. ACTS Information Center, <http://acts.nerisc.gov>
2. ACTS Tools web-page: <http://acts.nerisc.gov/tools>
3. E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*, Third Edition. SIAM: 1999.
4. Satish Balay, William Gropp, Lois Curfman McInnes, and Barry F. Smith. *PETSc users Manual*, ANL-95/11 –Revision 2.1.1, Argonne National Laboratory, 2001.
5. Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient Management of Parallelism in Object Oriented Numerical Software Libraries. In: E. Arge, A. M. Bruaset, and H. P. Langtangen Eds. *Modern Software Tools in Scientific Computing*. Birkhauser Press, 1997:163—202.
6. CCA-Forum, <http://www.cca-forum.org>
7. CCTTSS, <http://www.cca-forum.org/ccttss>
8. Chow, E., A.J. Cleary, and R.D. Falgout. Design of the Hypre Preconditioner Library. In: Mike Henderson, Chris Anderson, and Steve Lyons Eds. *Proc. of the SIAM Workshop on Object Oriented Methods for Inter-operable Scientific and Engineering Computing*. 1998.
9. L. A. Drummond, J. D. Farrara, M. Gupta, C. R. Mechoso, and J. A. Spahr. The UCLA Earth System Model: Computational Challenges. In: *Proceedings from 1999 International Conference on Mission Earth*. The Society for Computer Simulations International, San Francisco, California, January 17-20, 1999; 38-42.

10. I. Foster and C. Kesselman Eds. 1999. *The Grid, Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, Inc, San Francisco, CA.
11. Ian Foster. The Grid: A New Infrastructure for 21st Century Science, *Physics Today*. February, 2002.
12. FUN3D, <http://fmad-www.larc.nasa.gov/~wanderso/Fun>
13. G. A. Geist, J. A. Kohl, P. M. Papadopoulos. CUMULVS: Providing Fault-Tolerance, Visualization and Steering of Parallel Applications. *International Journal of High Performance Computing Applications* 1997; Volume 11, Number 3: 224-236.
14. W. Gropp, D. Kaushik, D. Keyes and B. Smith. Toward Realistic Performance Bounds for Implicit CFD Codes. *In Proceedings of Parallel CFD'99*. Elsevier. 1999.
15. LAPACK, <http://www.netlib.org/lapack>
16. PETSc-FUN3D, <http://www.mcs.anl.gov/petsc-fun3d>
17. C.R. Mechoso, J.D. Farrara, L.A. Drummond, J. Spahr and J.-Y. Yu. "An atmosphere-ocean model: Code optimization and application to El Nino". In: D. A. Gabriel Ibarra-Berastegi, C. A. Brebbia and P. Zanetti, Eds. *Development and Application of Computer Techniques to Environmental Studies*. WIT Press, 1999: 544.
18. Nieplocha, RJ Harrison, and RJ Littlefield. Global Arrays: A portable 'shared-memory' programming model for distributed memory computers. In *Proceedings from Supercomputing'94*. 1994:340-349
19. NWChem, <http://www.emsl.pnl.gov:2080/docs/nwchem>
20. ODE, <http://www.netlib.org/ode>
21. PETSc Website, <http://www.mcs.anl.gov/petsc>

22. W. Press, S. Teukolsky, W. Vetterling and B. Flannery. *Numerical Recipes in Fortran 77*. Cambridge University Press. 1992.
23. W. Press and S. Teukolsky. *Numerical Recipes: Does This Paradigm Have a Future?* June 6, 1997.
24. SciDAC/DOE, <http://www.er.doe.gov/scidac>
25. J. N. Shadid, R. S. Tuminaro. Iterative Methods for Nonsymmetric Systems on MIMD Machines. In: Dongarra, Messina, Sorensen, and Voight Eds. *Proceedings of the Fifth SIAM Conference on Parallel Processing for Scientific Applications*. Houston, TX, March 25-17, 1991.
26. S. Shende, A. D. Malony, J. Cuny, K. Lindlan, P. Beckman and S. Karmesin. Portable Profiling and Tracing for Parallel Scientific Applications using C++. In: *Proceedings of SPDT'98: ACM SIGMETRICS Symposium on Parallel and Distributed Tools*. 1998:134-145.
27. SLEPc, <http://acts.nersc.gov/workshop/slides/roman.pdf>
28. Trilinos, <http://www.cs.sandia.gov/Trilinos>
29. R. S. Tuminaro, M. Heroux, S. A. Hutchinson, and J. N. Shadid. *Official Aztec User's Guide: Version 2.1*. 1999. <http://www.cs.sandia.gov/CRF/aztec1.html>.